

Information Disclosure as a Means to Security

Zinovi Rabinovich
Mobileye Vision Technologies
Ltd
Jerusalem, Israel
zr@zinovi.net

Albert Xin Jiang
Trinity University
San Antonio, USA
xjiang@trinity.edu

Manish Jain
Armorway
Los Angeles, USA
manish@armorway.com

Haifeng Xu
University of Southern
California
Los Angeles, USA
haifengx@usc.edu

ABSTRACT

In this paper we present a novel Stackelberg-type model of security domains: *Security Assets Assignment with Information Disclosure (SASI)*. The model combines both the features of the Stackelberg Security Games (SSGs) model and of the Bayesian Persuasion (BP) model. More specifically, SASI includes: a) an uncontrolled, exogenous security state that serves as the Defender's private information; b) multiple security assets with non-accumulating, target-local defence capability; c) a pro-active, verifiable and public, unidirectional information disclosure channel from the Defender to the Attacker. We show that SASI with a non-degenerate information disclosure can be arbitrarily more efficient, than a "silent" Stackelberg assets allocation. We also provide a linear program reformulation of SASI that can be solved in polynomial time in SASI parameters. Furthermore, we show that it is possible to remove one of SASI parameters and, rather than require it as an input, recover it by computation. As a result, SASI becomes highly scalable.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Intelligent Agents

General Terms

Algorithms, Security

Keywords

security games; information disclosure

1. INTRODUCTION

In recent years the issue of security asset assignment has been given ever increasing attention. In particular, Stackelberg Security Games (SSGs) between a Defender and an Attacker have gained popularity as a rigorous theoretical and a successful practical solution. Deployed solutions range from LAX Air Port and Federal Air Marshals Service [12] to US Coast Guard patrols [22, 10].

Most of these works assume that the attacker observes the defender's history of daily security allocations, thus learning the de-

fender's (possibly mixed) strategy, before deciding on a target to attack. The actual security allocation on the day of the attack is not known by the attacker. The Attacker's powers of observation are accepted as an uncontrollable given. While some works have investigated the impact of assuming a bound on these abilities (e.g. [2, 17, 5]), the main assumption holds: there is no *pro-active* control over the observability of security measures by the attacker. In real life, however, this is not a fixed rule.

In this paper we try to answer the question: when is it advantageous for the defender to reveal additional private information to the attacker? Such private information can include the defender's security allocation for the day, as well as the defender's knowledge about targets' values and/or vulnerabilities. For example, in Section 2 we describe a domain that is inspired by current efforts on wildlife protection [26], in which the defender has private information about the location of the wild animals.

Before introducing our model, we first briefly discuss relevant literature. Much debate has been devoted to pros and cons of security visibility. E.g. Powell's model [18] shows that hiding a target's vulnerability may actually be more important than defending it in case of a factual attack. On the other hand, as Zhuang and Bier [27] point out in their comparative review on secrecy and deception in homeland-security, there are also many examples where disclosure of the available security is beneficial (e.g. Bier et al [4]).

Furthermore, multiple studies on multi-agent domains have proactively exploited the asymmetry in knowledge and actuation abilities. In fact, the controlled information disclosure consistently benefits the knowledgeable agent over the actuating agent, ranging from concerns for social impact (see e.g. Emek et al [9], Guo and Deligkas [11], Dughmi et al [8]) to individual human interactions (see e.g. Schweizer and Szech [19], Azaria et al [3]) to studies of factual terror activities (see e.g. Serra and Subrahmanian [20, 21]).

While there are many models that capture strategic information disclosure, the most relevant model for our purpose is Kamenica and Gentzkow's [14] *Bayesian Persuasion (BP)*. In more detail, the knowledgeable party (termed the *Sender*) is the only one who possesses the relevant, specific knowledge that affects the utility of both participants, while the other player (the *Receiver*) is the only one who can act in order to bring about those rewards. The disbalance of power forces the Sender to provide only verifiable and *usable* information, otherwise it would be discarded and have no effect. In other words, the utility of the Receiver has to rise as the result of adopting the information provided by the Sender. Of course, there's no better lie than a half-truth, and the Sender does just that – provide only partial, probabilistic information that

creates a bias in the Receiver’s knowledge which, in turn, would bias its actions and benefit the Sender. On the other hand, unlike the BP model, in security domains the knowledgeable party (defender) also need to decide on (randomized) security allocations.

It is interesting to note that persuasion processes and real world Defender-Attacker interactions are time consuming. E.g., Agmon [1] considers the effects of an Attacker interception timing on patrol strategies. Her work would address such an issue as a nature reserve patrol, where stopping a poacher in the beginning of a hunt is better than catching him after the fact. The time extent of a Defender-Attacker interaction produces even further difficulties and challenges, if one considers the sequence in which observations, decisions and choices are made. In particular, [25] consider the situation where the target choice occurs before any additional local observations or the attack execution decision take place. They demonstrate the failure of the standard SSG model to handle such a separation in a natural security domain, underlining the importance of timing, and provide an effective alternative model.

Nonetheless, the common assumption, adopted here, is that an Attacker gathers all the information prior to target selection, and the attack is instantaneous afterwards. This would be justified, among other reasons, by the need to maintain communication silence during attack execution and the inability to maintain secrecy for prolonged periods of time, once the attack target and means have been chosen.

Against this background, we present here a new model, *Security Assets aSsignment with Information disclosure (SASI)*, that combines the features of the SSGs and BP. More specifically, SASI includes: a) an uncontrolled, exogenous security state that serves as the Defender’s private information; b) multiple security assets with non-accumulating, target-local defence capability; c) a pro-active, verifiable and public, uni-directional information disclosure channel from the Defender to the Attacker. We show that SASI with a non-degenerate information disclosure can be arbitrarily more efficient, than a “silent” Stackelberg assets allocation. We also provide a linear program reformulation of SASI that can be solved in time polynomial in SASI parameters. Furthermore, we show that it is possible to remove one of SASI parameters and, rather than require it as an input, recover it by computation. As a result, SASI becomes highly scalable.

2. EXAMPLE DOMAIN: RHINO’S DAY

In a wildlife sanctuary, there are two types of pastures: muddy low-lands and mountain slopes. Due to the terrain, there are two pastures of each type: “The Bog” and “The Mire” in the low-lands, and “The Hill” and “The Flat Top” in the mountains. Two very rare white rhinos, Bertha and Bob, have been radio-tagged, vaccinated and released in the sanctuary. Since then they keep randomly visiting all four pastures for their daily grazing. For some reason Bertha and Bob always choose the same type of pasture, yet always graze separately, never visiting the same pasture together. Even though their choice between the mountain slopes and the muddy low-lands appears completely random, Bertha and Bob tend to frequent the latter a bit more for the added benefit of mud bathing.

Unfortunately, Bertha and Bob’s life is not idyllic. Alice, the poacher, is always trying to sneak at them during the day, while they graze, and kill them. To protect the endangered animals, Debby, the ranger, has set up guard posts on each of the four pastures. Having consulted the radio-tag data in the morning, she mans one of the posts for the day. If caught trying to kill Bertha or Bob, Alice will be arrested by Debby and fined. If Debby fails to protect the rhinos, the repeated introduction of the species to the sanctuary will be very costly.

Since the sanctuary is government controlled, all the scientific data and the employee performance is a matter of public record and knowledge. In fact, to demonstrate the government’s battle with corruption, all Debby’s activity is logged and can be verified after the fact. Furthermore, to promote the sanctuary, Debby has to give regular and frequent public interviews. Alice, therefore, is fully aware of rhino habits and any protective measures and protocols that Debby establishes.

As a result, a constant scheduling for guard posts can not be used. Alice would simply study the sanctuary activity reports and anti-synchronise her attacks with Debby’s guard duties. Hence, Debby has to deploy and commit to a randomised scheduling.

At this point, the domain and the difficulties faced by Debby appear to be just a standard security scheduling domain. Much like scheduling Air Marshals (e.g. [23, 12]). However, there is an additional domain feature that Debby can employ, and which places is outside the scope of standard security scheduling models, such as SSGs.

As we have mentioned, Debby is forced to give regular *public* interviews. In effect, this is a uni-directional communication channel between her and Alice. As we will show, besides using randomised guard scheduling, Debby can also release additional information to “warn off” Alice’s attacks.

The intuition is as follows. While directly releasing information about her specific deployment is counterproductive, Debby can release information that would speak about the correlation between her choice of the guard post and the rhino’s pasture choice. Effectively, this will create an impression on Alice that rhinos are far more likely to be protected, than not. Practically, this is achieved by randomising the content of Debby’s public interviews. Since it is also logged and verifiable after the fact, Debby can construct and commit to a specific conditional distribution of public utterances. This distribution will serve as a persuasive communication, similarly to the one used by Bayesian Persuasion[14], from Debby to Alice.

In fact, as we will show in Section 5, Debby can indeed achieve higher deterrence by speaking in public. However, to do so formally we first must introduce a model capable of fully capturing the Rhino’s Day domain. We do so in the next two sections.

3. NOTATION

We will use the following notation and symbols. The set of all integers ranging from 1 to an upper bound L will be denoted by $[L]$. Same lower case letter will denote an element of the set, $l \in [L]$. In turn, the space of all discrete distributions over a finite set $[L]$ will be denoted by $\Delta([L])$.

We will use Greek letters to denote a specific distribution, and adopt the functional notation. In particular, a conditional distribution will be seen as function from the conditioning set into a space of distributions, e.g. $\rho : [S] \rightarrow \Delta([L])$. However, to underline the conditioning we will use $\rho(\cdot|s)$ to denote function’s value at specific s , and $\rho(l|s)$ to denote the conditional probability of a specific event.

We will use capital calligraphic letters to denote matrices, e.g. \mathcal{B} , and the functional form $\mathcal{B}(l, s)$ to denote the element at row l and column s .

A superscript will always denote the agent in control of (or characterised by) the term, and a subscript denote a feature of that term. For example, consider the term t_*^A , it is an element of the discrete set $[T]$ of targets, where the superscript A denotes that it is controlled by the Attacker agent, and the subscript $*$ that it is optimal in some sense.

4. SASI: MODEL

In this section we will describe our model of *Security Assets as-Signment with Information disclosure (SASI)* between a Defender and an Attacker agents. We will begin by describing the features, parameters and assumptions of the interaction, and proceed to formally define the optimisation each agent needs to perform to maximise its utility. Overall, the model is a fusion of the Stackelberg Security Game (SSG) and the Bayesian Persuasion (BP) models, and at the end of this section we describe the parameterisations under which SASI degenerates into SSG and BP.

4.1 Interaction Features, Parameters and Assumptions

The Defender is tasked with the allocation of M security assets to protect T targets (with $M < T$). Every asset can protect only the target it is directly assigned to, and no increased protection is provided by assigning more than one asset to a target. We assume that the Defender has L complete, non-redundant asset allocations or **plays**, which we organise into a binary **playbook** matrix \mathcal{B} of size $L \times T$. $\mathcal{B}(l, t) \in \{0, 1\}$ denotes whether there is a security asset assigned to target t by the play l , and for all $l \in [L]$ holds $\sum_{t \in [T]} \mathcal{B}(l, t) = M$. In turn, the Attacker is tasked with choosing a single target that it will attempt to damage.

However, the utility gained by each agent is dictated by three factors: the Defender's play, the Attacker's target of choice, and the overall security and vulnerability **state**. We assume that there is a finite set of S such states, hence the utility function for the Defender has form: $u^D : [L] \times [T] \times [S] \rightarrow R$; and for the Attacker: $u^A : [L] \times [T] \times [S] \rightarrow R$. For any specific interaction between the Defender and the Attacker, the security state is sampled from a publicly known distribution $\lambda \in \Delta([S])$. Yet, the specific sampled instance of the state is only known to the Defender.

In essence, the overall security state is the Defender's private information (a.k.a. Harsanyi type). In particular, we will allow the Defender to deploy a conditional mixed strategy of the form $\rho : [S] \rightarrow \Delta([L])$, where $\rho(l|s)$ is the probability of using play l given that the overall security state is s . Consequently, both agents will seek to optimise their expected utility with respect to the strategy randomisation and the general security state uncertainty.

As with the classical SSGs, we assume that utility functions are decomposable, based on the availability of a security asset at the attacked target. More specifically, let $u_o^D : [T] \times [S] \rightarrow R$ (respectively u_o^A) denote the utility of the Defender (Attacker) from an unprotected, open target being attacked, while $u_c^D : [T] \times [S] \rightarrow R$ (respectively u_c^A) denotes the utility from an attack on a protected, covered target. Then:

$$\begin{aligned} u^D(l, t, s) &= \mathcal{B}(l, t)u_c^D(t, s) + (1 - \mathcal{B}(l, t))u_o^D(t, s), \\ u^A(l, t, s) &= \mathcal{B}(l, t)u_c^A(t, s) + (1 - \mathcal{B}(l, t))u_o^A(t, s). \end{aligned}$$

On the other hand, deviating from SSGs, we allow a pro-active, public, uni-directional communication from the Defender to the Attacker. Specifically, we allow the Defender to select a set of messages of size I , and devise an information disclosure rule $\pi : [L] \times [S] \rightarrow \Delta([I])$. In particular, $\pi(i|l, s)$ is the probability of sending to the Attacker the message i , given that the play l has been selected under the overall security state s . To build intuition, think of a battle cry, or the call of the Quarterback in the American Football.

Finally, we will adopt the Stackelberg assumption, i.e. the Attacker is fully aware of the strategy and the disclosure rule adopted (and committed to) by the Defender *before* selecting a target to attack. In particular, the Attacker can use the Defender's message to reduce the uncertainty about the current security state and the

play selected by the Defender. We assume that the Attacker uses Bayesian inference to this end. As a result, the overall interaction between the Defender and the Attacker proceeds as follows:

- Interaction preliminaries:
 - The distribution $\lambda \in \Delta([S])$ is announced to both the Defender and the Attacker,
 - The Defender calculates, announces and commits to the play strategy $\rho : [S] \rightarrow [L]$ and the information disclosure rule $\pi : [L] \times [S] \rightarrow \Delta([I])$, including the message set $[I]$.
- Interaction instance:
 - A specific instance s is sampled, $s \sim \lambda$, and privately given to the Defender,
 - The Defender samples $l \sim \rho(\cdot|s)$, and then $i \sim \pi(\cdot|l, s)$,
 - The message i is announced,
 - Based on the received message, and the knowledge of π , ρ and λ , the Attacker selects the target, $t_*(i)$, to attack,
- Interaction conclusion:
 - The Defender receives utility $u^D(l, t_*(i), s)$
 - The Attacker receives utility $u^A(l, t_*(i), s)$

4.2 Agents' Expected Utility and Optimal Choice

Let us now formalise the manner in which the Defender and the Attacker optimise their choices.

The Attacker's utility depends on three parameters, but it controls only one of them: the attack's target. The Defender's play and the security state are known to the Attacker only in the form of a distribution defined by ρ and λ . The Attacker has to average over these **prior beliefs** to obtain the utility's expectation:

$$u^A(t) = \sum_{l, s} \rho(l|s)\lambda(s)u^A(l, t, s)$$

However, once the Defender's message is announced, the Attacker can update its beliefs using Bayesian inference to obtain the utility expectation conditioned on the message:

$$u^A(t, i) = \frac{1}{Z(i)} \sum_{l, s} \pi(i|l, s)\rho(l|s)\lambda(s)u^A(l, t, s),$$

$$\text{where } Z(i) = \sum_{l, s} \pi(i|l, s)\rho(l|s)\lambda(s).$$

As a result, assuming that the Attacker is a rational player, the optimal target given the prior beliefs is $t_* = \arg \max_t u^A(t)$. In turn, the optimal choice after receiving a particular message becomes:

$$t_*(i) = \arg \max_t u^A(t, i). \quad (1)$$

Notice that the utility expectation, and the optimal target choice t_* and $t_*(i)$, actually depend on π and ρ . Although formally this necessitates the functional notation, e.g. $t_*(i, \pi, \rho)$, we will omit π and ρ for brevity, where they are clear from context, as we have done above.

Now, having determined the expected utility and the optimal choice of the Attacker, we can calculate the expected utility of the Defender from a particular strategy, ρ , of selecting a play and a disclosure rule, π :

$$u^D(\pi, \rho) = \sum_{i, l, s} \pi(i|l, s)\rho(l|s)\lambda(s)u^D(l, t_*(i), s) \quad (2)$$

Therefore, the optimal choice for the Defender is determined by the optimisation problem:

$$(\pi_*, \rho_*) = \arg \max_{\pi, \rho} u^D(\pi, \rho) \quad (3)$$

4.3 Relationship to other models

Notice that if $L = 1$, then independently of the number of targets and security resources, the only strategic depth that the Defender has is expressed by the information disclosure rule π . As a result, with $L = 1$, SASI degenerates and becomes equivalent to Bayesian Persuasion.

Obviously, setting $I = 1$, i.e. limiting the Defender to a single utterance, converts the SASI model into an SSG. The Defender simply has no strategic ability beyond security resource allocation.

Interestingly, however, the same occurs when setting $S = 1$. But here the reason is a bit more intricate. Because there is no additional variability in the utility of an attack due to the hidden security state, any additional information revealed by the Defender will disclose a part of the security resources allocation. It turns out that the optimal solution in this case coincides with the optimal SSG solution, i.e., it is optimal to set $I = 1$ and reveal no information at all. We will derive this in a more formal manner in Corollary 2.

If the game is zero sum, i.e., $u^D(l, t, s) + u^A(l, t, s) = 0$ for all l, t, s , then it is optimal to reveal no information, in which case the optimal solution coincides with the SSG solution and the maxmin solution. To see this, let v be the attacker’s expected utility in the maxmin solution. Now consider an information disclosure scheme with $I > 1$. Because the game is zero-sum, the attacker can always guarantee at least v utility by ignoring the revealed information and playing his minmax strategy. So the defender cannot get better than $-v$ utility.

To further demonstrate the properties of SASI, in the following section we will present a formal model of our Rhino’s Day example.

5. SASI RHINO’S DAY

Let us have a more formal look at the Rhino’s Day domain. There is only one security asset to be assigned (Debby herself), $M = 1$. There are five targets, $M \ll T = 5$. Targets 1(one) through 4(four) stand for the pasture (“The Hill”, “The Flat Top”, “The Bog” and “The Mire”), while the fifth “dummy” target denotes Alice’s home. The natural playbook for Debby has four plays, $L = 4$, each stationing her at a guard post of a particular pasture, so that $\mathcal{B}(l, t) = 1 \iff l = t$. There are two general security states, $S = 2$, to denote the choice of the mountain slopes ($s = 1$) and muddy low-lands ($s = 2$) pasture types. Since Bertha and Bob graze almost randomly, but tend towards the low-lands, $\lambda(s = 1) = 0.5 - \epsilon$ and $\lambda(s = 2) = 0.5 + \epsilon$.

The utility functions $u_c^D, u_o^D, u_c^A, u_o^A$ clearly depend on three parameters, the combination of the play (Debby’s location), the selected target (Alice’s location) and the security state (pasture type selected by Bertha and Bob). However, it will be easier to summarise them (see Table 1) by grouping these combinations based on the presence of Bertha and/or Bob at an attacked pasture, and the set of combinations where Alice is staying home. E.g. for a situation, (l, t, s) , where either Bob or Bertha are present at the attacked pasture, has to hold that $t = 2 * s - 1$ or $t = 2 * s$. At the same time, if neither Bertha nor Bob are present, then $t \neq 2 * s - 1$ and $t \neq 2 * s$. Finally, for all (l, t, s) where Alice stays home, $t = 5$.

Now, the requirement that Debby provides Alice with useful, though not necessarily complete information, translates into a set of ambiguous utterances, $[I]$, where the ambiguity is expressed via

Situation class	u_c^D	u_o^D	u_c^A	u_o^A
Bertha/Bob is present	0	-10	-1	2
Bertha/Bob are absent	0	0	-0.5	-0.5
Alice at home	0	0	0	0

Table 1: Rhino’s Day: Basic payoffs table

the information disclosure rule $\pi : [L] \times [S] \rightarrow \Delta([I])$. In other words, Alice knows how likely Debby is to say something, given the selected pasture type and which guard post Debby will actually man.

Our example domain also clearly demonstrates the connection with SSGs and BPs. If Debby is prevented from public speaking ($I = 1$), the domain immediately becomes an SSG. On the other hand, if Debby is released from her guard duties and becomes the spokesperson, i.e. is limited only to public speaking without actual protection, then she has only persuasion at her disposal, and the domain turns into BPs.

More importantly, however, Rhino’s Day domain serves as in a key tool to demonstrate SASI effectiveness. Specifically, that the coefficient of utility boost provided by a SASI solution is not bounded within the security games class. The next section contains a formal statement to this effect, and its proof by construction relies on the detailed computation of Debby’s and Alice’s utility gains in the Rhino’s domain.

6. SASI: BENEFITS AND COMPLEXITY

In this section we provide our results regarding the efficiency and the computational complexity of a SASI model.

It is easy to see that a SASI solution is at least as efficient as an SSG solution, simply because SSG equilibrium is a solution of SASI with unit sized messages set. However, a true SASI equilibrium, one that employs both resource allocation *and* information disclosure can be far more efficient.

Theorem 1. *For any N , there is a SASI instance given by a tuple $\langle M, T, L, S, \lambda, u_o^D, u_o^A, u_c^D, u_c^A \rangle$ with non-positive Defender utilities so that $\frac{u^D(\pi_1, \rho)}{u^D(\pi_I, \rho)} > N$, where (π_1, ρ) and (π_I, ρ) are optimal security assets allocation strategies with disclosure rules of size 1(one) and $I > 1$ respectively.*

PROOF. Consider the Rhino’s Day domain formulated as SASI. Notice that Alice is indifferent between attacking Bob or Bertha. Intuitively, if she chooses to attack, the only way to maximise the likelihood of catching her is for Debby to protect either Bob or Bertha with probability 0.5. In other words, the optimal assets allocations strategy is given by:

$$\rho(l|s) = \begin{cases} 0.5 & l = 2 * s - 1 \vee l = 2 * s \\ 0 & \text{otherwise} \end{cases}$$

Since it is more likely for Bob and Bertha to visit a muddy low-lands pasture $\lambda(s = 1) = 0.5 + \epsilon$, then, if she decides to attack and without any further information, Alice will stake either one of $t = 3$ or $t = 4$ (“The Bog” or “The Mire”). If Bob and Bertha indeed decide on a low-lands pasture, Alice will then be caught with probability 0.5, otherwise she’ll just be wasting her time since all the action will be elsewhere. As a result: $u^A(t = 3) = u^A(t = 4) = (0.5 * 2 + 0.5 * (-1)) * (0.5 + \epsilon) + (-0.5) * (0.5 - \epsilon) = \epsilon$, $u^A(t = 1) = u^A(t = 2) = -\epsilon$ and $u^A(t = 5) = 0$. Hence, Alice would indeed prefer to execute an attack and choose “The Bog”. Debby’s utility will then become $u^D(\pi_1, \rho) = -5 * (0.5 + \epsilon) = -2.5 - 5\epsilon$.

For the sake of developing the necessary computational intuition, let us show the aforementioned calculation of $u^A(\cdot)$, $\rho(l|s)$ and $u^D(\pi_1, \rho)$ more explicitly. Afterwards we will omit explicit algebraic computations for brevity.

Recall, $u^A(t) = \sum_{l,s} \rho(l|s)\lambda(s)u^A(l, t, s)$, and $\lambda(s=1) = 0.5 - \epsilon$, $\lambda(s=2) = 0.5 + \epsilon$.

Notice that for $t = 5$, $u^A(l, t = 5, s) = 0$ for all $l \in [L]$, $s \in [S]$, so that $u^A(t = 5) = 0$. I.e. the utility of Alice staying home is constant zero for any policy that Debby may adopt.

Let's now take a closer look at $u^A(t = 1), \dots, u^A(t = 4)$. Substituting the known values into the formula of $u^A(t)$ we obtain:

$$u^A(t = 1) = -3\rho(l = 1|s = 1) * (0.5 - \epsilon) + 0.75 - 2.5\epsilon \quad (4)$$

$$u^A(t = 2) = -3\rho(l = 2|s = 1) * (0.5 - \epsilon) + 0.75 - 2.5\epsilon \quad (5)$$

$$u^A(t = 3) = -3\rho(l = 3|s = 2) * (0.5 + \epsilon) + 0.75 + 2.5\epsilon \quad (6)$$

$$u^A(t = 4) = -3\rho(l = 4|s = 2) * (0.5 + \epsilon) + 0.75 + 2.5\epsilon \quad (7)$$

As a result, if $\rho(l = 1|s = 1) > \rho(l = 2|s = 1)$, then Alice would rather attack target $t = 2$ ("The Flat Top"), and otherwise she would prefer to attack $t = 1$ ("The Hill"). Similarly, if $\rho(l = 3|s = 2) > \rho(l = 4|s = 2)$, then Alice prefers to attack $t = 4$ ("The Mire") and otherwise $t = 3$ ("The Bog"). In essence, Alice will always prefer the target opposite to the comparison of $\rho(l = 2*s - 1|s)$ and $\rho(l = 2*s|s)$. W.l.g. assume that $\rho(l = 2*s - 1|s) > \rho(l = 2*s|s)$.

Taking a closer look at $u^D(\pi_1, \rho)$ it is clear that as $\rho(l = t_*|s)$ grows so does the utility of the Defender Debby. Let's assume that Alice has chosen $t_* = 1$. Then it had to hold that $\rho(l = 2|s = 1) > \rho(l = 1|s = 1)$. The maximal value of $\rho(l = 1|s = 1)$ under these conditions is 0.5. In fact, for any choice of t_* this holds, i.e. $\rho(l = 2*s - 1|s) = \rho(l = 2*s|s) = 0.5$. Substituting these values into Equations 4-7 we obtain $u^A(t = 1) = u^A(t = 2) = -\epsilon$, $u^A(t = 3) = u^A(t = 4) = \epsilon$, and $u^A(t = 5) = 0$. Hence, Alice has to choose $t_* = 3$ ("The Bog"), and further substitution yields $u^D(\pi_1, \rho) = -2.5 - 5\epsilon$ as required.

Now, consider, on the other hand, an information disclosure scheme with two messages, $I = 2$, where for any $l \in [L]$ we define $\pi_I(i = 0|s = 0, l) = 1$, $\pi_I(i = 0|s = 1, l) = \frac{0.5 - \epsilon}{0.5 + \epsilon}$ and $\pi_I(i = 1|s = 1, l) = \frac{2\epsilon}{0.5 + \epsilon}$. Upon receiving signal $i = 0$, Alice infers a posterior distribution on l and s , with the following properties. If message $i = 0$ is received then the posterior likelihood of $s = 0$ and $s = 1$ will be equal. As a result $u^A(t, i = 0) = 0$ for any target. Since we assume tie breaking in favour of the Defender, Alice will simply choose $t = 4$, i.e. stay home. Debby's utility in this case will also be zero. On the other hand, if $i = 1$ is received (which happens with probability of 2ϵ), then the likelihood of $s = 1$ is one, i.e. Alice knows for sure that Bob and Bertha have chosen the muddy low-lands. In particular, $u^A(t = 2, i = 1) = u^A(t = 3, i = 1) = 0.5$ and Debby's utility will be -5 . Since message $i = 0$ is sent with probability $1 - 2\epsilon$ and the message $i = 1$ is sent with probability 2ϵ , the overall utility of the Defender Debby will be $u^D(\pi_I, \rho) = 0 * (1 - 2\epsilon) - 5 * 2\epsilon = -10\epsilon$.

Notice that $\lim_{\epsilon \rightarrow 0} \frac{u^D(\pi_1, \rho)}{u^D(\pi_I, \rho)} = \frac{-2.5 - 0.5\epsilon}{-10\epsilon} = \infty$, hence for any N we can choose ϵ so that $\frac{u^D(\pi_1, \rho)}{u^D(\pi_I, \rho)} > N$. In other words, without information disclosure the Defender would lose at least N times more. For instance, when Rhino's are visiting the low-lands 10% more frequently, we'll have $\epsilon = 0.05$ and SASI will be able to gain more than 2.5 better utility than the standard SSG solution. Essentially this translates into more than double the security level for this rare animal breed. \square

Theorem 1 shows that there are cases where, compared to the SSG solution, SASI can guarantee an arbitrarily greater deterrence. However, it does not speak to how extensive the vocabulary, $[I]$, should be. As it turns out, in spite of SASI being an extension to Bayesian Persuasion, the bound devised by Kamenica and Gentzkow[14] remains effective. In fact, the following corollary shows this by following the same reasoning as the original Kamenica and Gentzkow[14] proof.

Corollary 1. *A SASI can be solved using a disclosure rule π that is based on no more than T messages.*

PROOF. Let $i_1 \neq i_2$ be two messages of an optimal solution pair $(\widehat{\pi}_*, \widehat{\rho}_*)$, so that $t_*(i_1|\widehat{\pi}_*, \widehat{\rho}_*) = t_*(i_2|\widehat{\pi}_*, \widehat{\rho}_*)$. We will show that an alternative solution $(\pi, \widehat{\rho}_*)$ can be devised, yielding the same utility to the Defender, where messages i_1 and i_2 are replaced by a single message \widehat{i} . Hence yielding the Corollary's bound.

Recall, $t_*(i, \pi, \widehat{\rho}_*) = \arg \max_{t \in [T]} u^A(t, i)$. We can rewrite this maximisation as a set of inequalities, reducing the normalisation factor $Z(i)$:

$$\forall t, \sum_{s,l} \pi(i|l, s) \widehat{\rho}_*(l|s) \lambda(s) \left[u^A(l, t_*, s) - u^A(l, t, s) \right] \geq 0 \quad (8)$$

A target t_* would only be the target with maximum utility if and only if it satisfies all inequalities above.

Let us now replace the two messages i_1 and i_2 with a single message \widehat{i} , and define $\pi(i|l, s) = \widehat{\pi}(i|l, s)$ for all $i \neq \widehat{i}$ and $\pi(\widehat{i}|l, s) = \widehat{\pi}_*(i_1|l, s) + \widehat{\pi}_*(i_2|l, s)$. It is easy to see that if Inequalities 8 regarding i_1 and i_2 under $(\widehat{\pi}_*, \widehat{\rho}_*)$ will hold, then so would the corresponding inequality for \widehat{i} under $(\pi, \widehat{\rho}_*)$. Furthermore, other messages will be unaffected. Obviously, π is well defined as conditional distribution. Similarly, the Defender's utility from adopting $(\pi, \widehat{\rho}_*)$ instead of $(\widehat{\pi}_*, \widehat{\rho}_*)$ will not change.

Thus we have obtained, another optimal solution to the given SASI with fewer messages. \square

Now, we have seen that the optimal solution is not too large, and it is enough to consider disclosure rules that have at most $I = T$ messages. Though a positive feature, by itself it does not necessarily guarantee that an optimal solution can be computed efficiently. Nonetheless, with the help of a few key observations, we construct a *Linear Program (LP)* that characterises the optimal solution. Thus, we show the following theorem.

Theorem 2. *SASI strategies optimisation, given by Equations 3 and 1, can be solved in time polynomial in T, L, S .*

Observation 1. *The proof of Corollary 1 implies that, in effect, every message can be interpreted as a direct advice, or order, to the Attacker to choose a specific target.*

For instance, in the case of the Rhino's Day domain, this meant that the message $i = 0$ translated directly into $t = 4$ (stay home), while $i = 1$ mapped into $t = 3$ ("The Mire").

Observation 2. *Computation of the optimal response t_* for any (π, ρ) is equivalent to a satisfiability check of T inequalities over STL variables.*

PROOF THEOREM2. Initially it seems that Equation 3 is an optimisation of a quadratic function due to $\pi(i|l, s)\rho(l|s)$ terms of $u^D(\pi, \rho)$, where $\pi(i|l, s)$ and $\rho(l|s)$ are variables of optimisation. However, because we optimise over π and ρ simultaneously, we can replace them with a single distribution $\xi : [S] \rightarrow \Delta([I] \times [L])$, and find π and ρ by marginalisation and conditioning. Rewriting

$u^A(t, i)$ and $u^D(\pi, \rho)$ in terms of ξ , and substituting into Equations 2 and 8, we recast the optimisation of Equation 3 as:

$$\begin{aligned} \xi_* &= \arg \max_{\xi} u^D(\xi) \quad s.t. & (9) \\ \forall i, l, s \quad 0 &\leq \xi(i, l|s) \leq 1 \\ \forall s \quad \sum_{i, l} \xi(i, l|s) &= 1 \\ \forall i, t \quad \sum_{l, s} \xi(i, l|s) \lambda(s) &\left[u^A(l, i, s) - u^A(l, t, s) \right] \geq 0 \end{aligned}$$

Notice that the above LP contains SLT number of variables, and $2 * SLT + S + T^2$ equations. Hence, is solvable in time polynomial in S, L, T . Marginalising $\rho_*(l|s) = \sum_i \xi_*(i, l|s)$ and conditioning $\pi_*(i|l, s) = \frac{\xi_*(i, l|s)}{\rho_*(l|s)}$ are also polynomial operations in S, L, T . \square

The proof of Theorem2 can also formally demonstrate our previous intuitive claim about the relationship between SSG and SASI with the degenerate number of security states, i.e. $S = 1$.

Corollary 2. *If $S = 1$ the optimal solution to SASI is equivalent to an SSG solution. I.e. no information is revealed by π_* .*

PROOF. With $S = 1$, the LP (9) degenerates to the LP formulation of Strong Stackelberg Equilibrium as described in [6]. \square

7. SCALING UP SASI

Thus far we have demonstrated that SASI can be more efficient than a regular SSG formulation for a domain, and that it can be solved in polynomial time in the number of security states, targets and plays. However, there is a caveat that has to be addressed to use SASI in practice: the playbook composition. Given a security domain with a non-trivial number of security assets $M \gg 1$ and a very large number of targets $T \gg M$. *A priori*, there are exponentially many possible non-redundant security resource allocations ($\binom{T}{M}$ possible plays in fact).

We circumvent this by the use of the *marginalisation trick*. The trick was used to resolve the very same scalability issue for SSGs: rather than calculating the optimal distribution over the complete set of assignments, SSG solvers calculate the *optimal marginal* probability of a target to be protected. Then by using the Birkhoff-von Neumann theorem, the relevant plays subset, i.e. the playbook, and the optimal mixture of plays is reconstructed. Notably, the size of such a playbook is polynomial in the number of targets and resources.

Now, for SASI the same marginalisation method can be applied, that is we can directly use marginal probabilities of coverage on individual targets in the LP for information revelation, and then afterwards extract playbook strategies from marginals using Birkhoff-von Neumann Theorem.

Specifically, instead of using $\xi(i, l|s)$, we establish new variables $\widehat{\xi}_0(i, t|s), \widehat{\xi}_1(i, t|s) \in [0, 1]$. In essence, $\widehat{\xi}_1(i, t|s)$ is the probability, given s , that message i was sent by the Defender and target t is covered; while $\widehat{\xi}_0(i, t|s)$ is the probability given s that i was sent by the Defender and target t is not covered. There are $T^2 S$ such variables. To be consistent with the existence of M security assets they have to satisfy $\forall s, i, \sum_t \widehat{\xi}_1(i, t|s) = M\pi(i|s)$. In other words, the total number of covered targets is equal to the number of security assets.

We can now reformulate the constraints and objectives of the LP that begins with the Equation 9. In particular, the Defender's

expected utility when the Attacker adopts the target advised by the received message is

$$u^D(\widehat{\xi}) = \sum_s \lambda(s) \sum_i \left[\widehat{\xi}_1(i, t|s) u_c^D(i, s) + \widehat{\xi}_0(i, t|s) u_o^D(i, s) \right].$$

To guarantee that the Attacker indeed adopts the advised target, it has to be optimal for him to do so. Recall that $u^A(t, i)$ is the utility of Attacker choosing t given message i , then it has to hold that $u^A(i, i) \geq u^A(t, i)$. We thus need to recast $u^A(t, i)$ in terms of $\widehat{\xi}$. To do so we will use a more detailed view of the Attacker's utility. Specifically, denote by $u^A(t, i, s)$ the expected utility from attacking target t after receiving message i in security state s . Then $u^A(t, i, s) = \frac{1}{Z(i)} (\widehat{\xi}_1(i, t|s) u_c^A(t, s) + \widehat{\xi}_0(i, t|s) u_o^A(t, s))$, and $u^A(t, i) = \sum_s \lambda(s) u^A(t, i, s)$. As before, rewriting the attacker's incentive constraint $u^A(i, i) \geq u^A(t, i)$ in terms of $\widehat{\xi}$ and multiplying by the normalization factor $Z(i)$, we have

$$\begin{aligned} \sum_s \lambda(s) \left(\widehat{\xi}_1(i, i|s) u_c^A(i, s) + \widehat{\xi}_0(i, i|s) u_o^A(i, s) \right) &\geq \\ \sum_s \lambda(s) \left(\widehat{\xi}_1(i, t|s) u_c^A(t, s) + \widehat{\xi}_0(i, t|s) u_o^A(t, s) \right) &\forall i, t \quad (10) \end{aligned}$$

Taking everything together, we have the following LP.

$$\begin{aligned} \arg \max_{\widehat{\xi}, \widehat{\pi}} u^D(\widehat{\xi}) \\ \widehat{\xi}_1(i, t|s), \widehat{\xi}_0(i, t|s) &\in [0, 1] \quad \forall i, t, s, \\ \widehat{\xi}_1(i, t|s) + \widehat{\xi}_0(i, t|s) &= \widehat{\pi}(i|s) \quad \forall i, t, s, \\ \sum_i \widehat{\pi}(i|s) &= 1 \quad \forall s, \\ \sum_t \widehat{\xi}_1(i, t|s) &= M\pi(i|s) \quad \forall i, s, \end{aligned}$$

Constraint (10).

The size of this LP is polynomial in T and S . Given a solution $\widehat{\xi}$, we would like to decompose it into a mixed strategy, i.e. a distribution $\xi(i, l|s)$, such that $\sum_l \xi(i, l|s) \mathcal{B}(l, t) = \widehat{\xi}_1(i, t|s)$ for all t, i, s .

This can be efficiently achieved by applying Birkhoff-vonNeumann for each i, s . Here it is useful to factor $\widehat{\xi}_1(i, t|s)$ as the product $\widehat{\rho}(t|i, s) \widehat{\pi}(i|s)$ where $\widehat{\rho}(t|i, s)$ is the probability of covering target t given i, s . Specifically, given each i, s , we can compute $\widehat{\rho}(t|i, s) = \widehat{\xi}_1(i, t|s) / \widehat{\pi}(i|s)$ for each target t . Then we are faced with the problem of coming up with a distribution $q(l|i, s)$ over some sub-set of pure resource assignments l (i.e. a playbook $\mathcal{B}_{i,s}$) that matches the marginal coverage $\widehat{\rho}(t|i, s)$ on targets, i.e., $\sum_l q(l|i, s) \mathcal{B}_{i,s}(l, t) = \widehat{\rho}(t|i, s) \forall t$. This can be solved efficiently using T support by Birkhoff-vonNeumann (see e.g., [15]). Notice that the optimal $\pi_*(i|l, s)$ and $\rho_*(l|s)$ can be either recovered from the above decomposition of $\widehat{\xi}_1(i, t|s)$, or recomputed directly with the union of the newly recovered playbooks $\cup_{i,s} \mathcal{B}_{i,s}$.

8. INFORMATION CONTENT OF SOLUTION

Before we conclude, we would like to present an information theoretic view of the SASI solution. More specifically, that for large problems SASI solution is equivalent in some sense to minimising the information theoretic error between the desired and implemented security event probabilities.

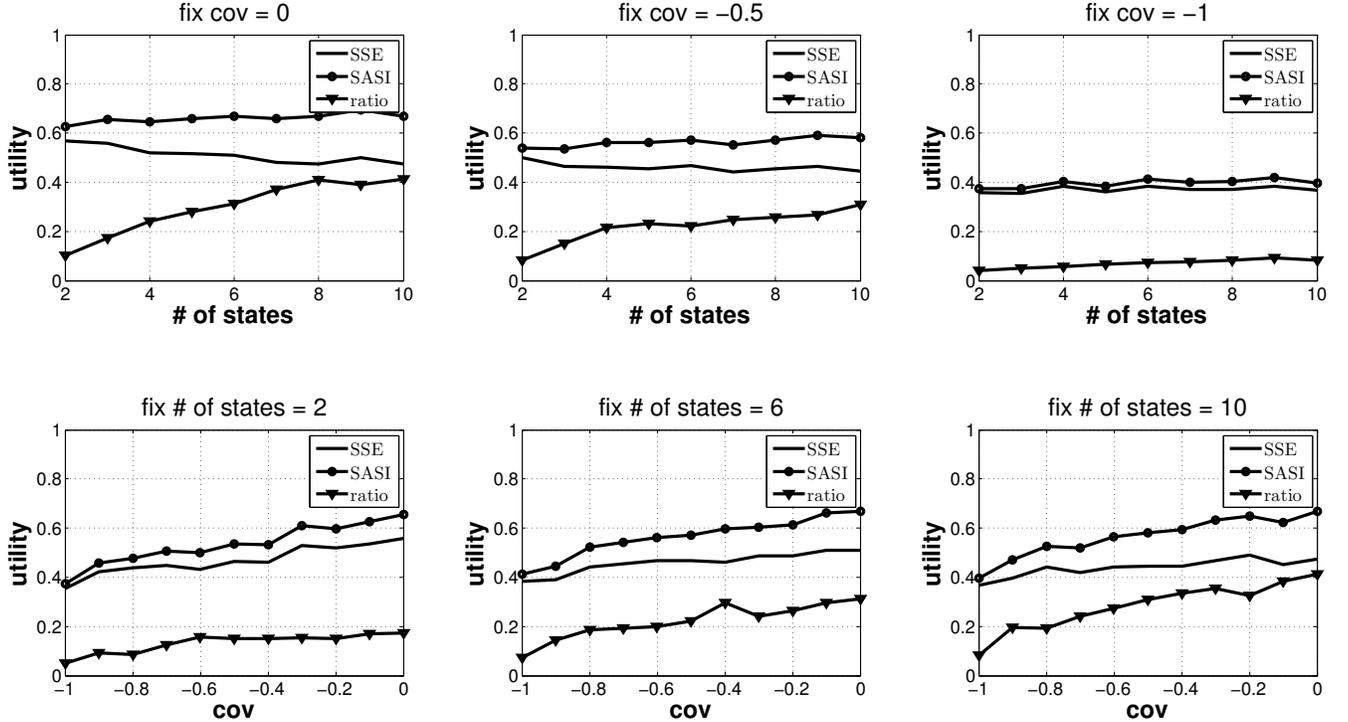


Figure 1: Utility comparison between SSE and SASI. The line "ratio" stands for the ratio $\frac{U_{SASI} - U_{SSE}}{U_{SSE}}$.

To show this formally, notice first that the optimal solution of Equation 9 is invariant to affine transformations of $u^D : [L] \times [T] \times [S] \rightarrow R$. In particular, we can assume that there is a distribution $\nu \in \Delta([L] \times [T] \times [S])$ so that $u^D(l, t, s) = \log \nu(l, t, s)$. In essence, $\nu(l, t, s)$ summarises both the tolerated and the required probability of security events. For instance, if $\mathcal{B}(l, t) = 0$, then $\nu(l, t, s)$ describes the relative tolerance towards an attack.

Now, let us denote the feasible set of SASI solutions, i.e. from which the Attacker benefits by following the advised targeting, by $\Sigma \subset \Delta([L] \times [T] \times [S])$. Formally:

$$\Sigma = \{\eta | \forall i, t \sum_{l,s} \eta(l, i, s) [u^A(l, i, s) - u^A(l, t, s)] \geq 0\}$$

Theorem 3. Assume that the disclosure rule has $I = T$ elements, and denote $\eta_*(l, i, s) = \pi_*(i|l, s)\rho_*(l|s)\lambda(s)$, where (π_*, ρ_*) is an optimal SASI solution. Then, as SLT increases, $D_{KL}(\eta_* || \nu) \rightarrow \min_{\eta \in \Sigma} D_{KL}(\eta || \nu)$.

In the theorem, D_{KL} stands for the Kullback-Leibler divergence measure between two distributions:

$$D_{KL}(\nu_1 || \nu_2) = \sum_i \nu_1(i) \log \frac{\nu_1(i)}{\nu_2(i)},$$

which, in information theory, is the formal measure of error when coding a signal that is distributed by ν_1 , using an approximate distribution ν_2 . Interpreting it from SASI point of view, it means that for larger problems, the optimal solution minimises the error between the real (implemented by (π_*, ρ_*)) and the tolerated (defined by u^D) security event probabilities. The proof of Theorem 3 is based on a combination of LP perturbed approximations introduced by Tsao and Fang [24] and D_{KL} geometry studied by Csizsar [7]. Unfortunately, space limitations force us to omit its details.

9. SIMULATIONS

In this section, we compare the utility of SSE and SASI on randomly generated security games. We generate the random payoffs for each target using the covariance random payoff generator [16]. Denote by $\mu[a, b]$ the uniform distribution on interval $[a, b]$, then we randomly generate the following random payoffs: $U_d^c \sim \mu[0, 0.5]$, $U_d^u \sim \mu[-0.5, 0]$, $U_a^c = aU_d^c + b\mu[-0.5, 0]$ and $U_a^u = aU_d^u + b\mu[0, 0.5]$, where $a = cov$, $b = \sqrt{1 - a^2}$. Here $cov \in [-1, 0]$ is the covariance parameter between defender's reward (or penalty) and attacker's penalty (or reward). So $cov = 0$ means a totally random payoff structure while $cov = -1$ means a constant-sum game. We further randomly generate L defender pure strategies, each of which is simply a subset of targets of size k , and $|S|$ randomly generated target permutations with a randomly generated distribution $p \in \Delta(S)$.

In all our simulations, we construct 15 targets and 10 pure strategies (i.e., $L = 10$) with 4 resources each (i.e., $k = 4$). Notice that the optimal defender mixed strategy of both SSE and SASI is invariant under linear shift, in order to compare the utility ratio, we further add each randomly generated target payoffs by 0.5, so that all the utilities are within interval $[0, 1]$. We explore the utility difference between SSE and SASI under different parameter pair $(|S|, cov)$. In particular, we take $|S|$ from 2, 3, ..., 10 and $cov = 0, -0.1, -0.2, \dots, -1$. For each parameter pair, we simulate 25 random games and average their utilities. Therefore, about $25 \times 11 \times 9 = 2475$ random security games are tested in total.

In all our simulations, SASI outperforms SSE which is as expected. Figure 1 provides part of the details of the utility comparison between SSE and SASI. In the top 3 figures, we fix the parameter cov and compare the utilities in terms of different $|S|$.

As the trend shows, SASI outperforms SSE more as $|S|$ increases. That means, as the states of nature grows, there is more “space” for the defender to manipulate the information to benefit herself. In the bottom 3 figures, we fix the parameter $|S|$ and compare the utilities in terms of different cov . The figures show that SASI outperforms SSE more as the defender and attacker’s payoffs have less covariance. This fits our intuition, since as the payoffs have less covariance, there are more cooperation “elements” of the game, which can be manipulated.

10. CONCLUSIONS AND FUTURE WORK

In this paper we develop a new model of interaction in security assets assignment domains. Specifically, while preserving the ability to generate a Stackelberg-optimal security assignment, we allow the Defender to use security information disclosure to bias the Attacker’s action, and further boost deterrence. We term the model *Security Assets Assignment with Information Disclosure (SASI)*. We show that SASI is solvable in time polynomial in the domain parameters. Furthermore, we demonstrate that the effects of the information disclosure, dictated by a SASI solution, can have arbitrarily large *positive* deterrence effect. We then show that it is possible to remove one of SASI parameters (the assignments playbook) and compose a computationally scalable solution. Based on this scalable formulation, we perform a set of experiments that further underline the effectiveness of SASI compared to that of SSE. In fact, our experimental data shows that SASI’s efficiency, relatively to that of SSE, increases with the number of targets.

Now, SASI makes an assumption that the probabilistic protocol of security assignment selection, as well as of the information disclosure, is a matter of public knowledge. We note that this naturally coincides with the transparency requirement in many for government institutions, such as the police force. Furthermore, even the more clandestine agencies operate under the assumption that their the Attacker has obtained the complete strategic information. This underlines the benefit of SASI, which exploits the Attacker’s knowledge as means of attack deterrence. Interestingly, the Attacker can not prevent the exploitation from occurring unless he is prepared to obtain sub-optimal utility. This is because SASI is *openly manipulative*. That is, the Attacker knows how the disclosed information is generated, he knows that it is generated to exploit him, but would still adopt the disclosed information, because the Attacker also knows that the information obtained is beneficial to his optimal decision making. After all, it would be better to withhold an attack, than getting caught.

Nonetheless, as a part of our future research we would like to provide additional guarantees for a *practical* implementation of our security assets assignment with information disclosure. In particular, we would like to extend our model to handle sequential decisions made by the Defender and the Attacker, including information accumulation, and temporal and spatial constraints. Most importantly, we would like to investigate the human response to the disclosed information.

Finally, it is worth mentioning two more extensions, kindly suggested by our reviewers. First, a question akin to “when to persuade” from Bayesian Persuasion can be posed. Our experiment statistics suggest that in at least half of all *general* security games, SASI will result in a utility boost greater than 40%. But can we describe a *sub-class* where such a guarantee always holds? Is there some game characteristic that would link with the degree of the utility boost? We have already made some progress in this direction, but a far more extended study is necessary.

Second, a form of *robust SASI* (similar to *robust SSE* by Jiang et al [13]) would be a natural next step. In more detail, it is necessary

to evaluate and support the stability of the information disclosure with respect to the precision of the Attacker model. Both with respect to the decision properties (e.g. human (ir)rationality) and the attacker’s utility (e.g. discrepancy between the Defender’s and the Attacker’s assignments of target importance).

REFERENCES

- [1] N. Agmon. On events in multi-robot patrol in adversarial environments. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, volume 2, pages 591–598, 2010.
- [2] N. Agmon, V. Sados, S. Kraus, and G. A. Kaminka. The impact of adversarial knowledge on adversarial planning in perimeter patrol. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 55–62, 2008.
- [3] A. Azaria, S. Kraus, C. V. Goldman, and Z. Rabinovich. Strategic information disclosure to people with multiple alternatives. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(4), 2015.
- [4] V. M. Bier, S. Oliveros, and L. Samuelson. Choosing what to protect: Strategic defensive allocation against an unknown attacker. *Journal of Public Economic Theory*, 9(4):563–587, 2007.
- [5] A. Blum, N. Haghtalab, and A. D. Procaccia. Lazy defenders are almost optimal against diligent attackers. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, pages 573–579, 2014.
- [6] V. Conitzer and D. Korzhyk. Commitment to correlated strategies. In *Proceedings of the 25th National Conference on Artificial Intelligence (AAAI)*, pages 632–637, 2011.
- [7] I. Csiszar. I-divergence geometry of probability distributions and minimisation problems. *The Annals of Probability*, 3(1):146–158, 1975.
- [8] S. Dughmi, N. Immorlica, and A. Roth. Constrained signalling in auction design. In *SODA*, 2014.
- [9] Y. Emek, M. Feldman, I. Gamzu, R. Paes Leme, and M. Tennenholtz. Signaling schemes for revenue maximization. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC ’12*, pages 514–531, New York, NY, USA, 2012. ACM.
- [10] F. Fang, A. X. Jiang, and M. Tambe. Optimal patrol strategy for protecting moving targets with multiple mobile resources. *JAIR*, 48:583–634, November 2013.
- [11] M. Guo and A. Deligkas. Revenue maximization via hiding item attributes. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI’13*, pages 157–163. AAAI Press, 2013.
- [12] M. Jain, J. Tsai, J. Pita, C. Kiekintveld, S. Rath, F. Ordonez, and M. Tambe. Software assistants for randomized patrol planning for the LAX airport police and the Federal Air Marshals service. *Interfaces*, 40(4):267–290, 2010.
- [13] A. X. Jiang, T. H. Nguyen, M. Tambe, and A. D. Procaccia. Monotonic maximin: A robust stackelberg solution against boundedly rational followers. In *Proceedings of the 4th Conference on Decision and Game Theory for Security (GameSec)*, pages 119–139, 2013.
- [14] E. Kamenica and M. Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- [15] J. Letchford and V. Conitzer. Solving security games on graphs via marginal probabilities. In *Proceedings of the 27th*

- AAAI Conference on Artificial Intelligence (AAAI-13), pages 591–597, 2013.
- [16] E. Nudelman, J. Wortman, Y. Shoham, and K. Leyton-Brown. Run the gamut: A comprehensive approach to evaluating game-theoretic algorithms. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 880–887. IEEE Computer Society, 2004.
- [17] J. Pita, M. Jain, M. Tambe, F. Ordonez, and S. Kraus. Robust solutions to Stackelberg games: Addressing bounded rationality and limited observations in human cognition. *Artificial Intelligence*, 2010.
- [18] R. Powell. Allocating defensive resources with private information about vulnerability. *American Political Science Review*, 101(4):799–809, 2007.
- [19] N. Schweizer and N. Szech. Optimal revelation life-changing information. Working Paper, 2012.
- [20] E. Serra and V. S. Subrahmanian. Should behavioral models of terrorist groups be disclosed? *IEEE Intelligent Systems*, 29(4), 2014.
- [21] E. Serra and V. S. Subrahmanian. A survey of quantitative models of terror group behavior and an analysis of strategic disclosure of behavior models. *IEEE Transactions on Computational Social Systems*, 1(1):66–88, 2014.
- [22] E. Shieh, B. An, R. Yang, M. Tambe, C. Baldwin, J. DiRenzo, B. Maule, and G. Meyer. Protect: A deployed game theoretic system to protect the ports of the united states. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012.
- [23] J. Tsai, S. Rathi, C. Kiekintveld, F. Ordonez, and M. Tambe. IRIS - a tool for strategic security allocation in transportation networks. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems – Industry Track*, pages 37–44, 2009.
- [24] H.-S. J. Tsao and S.-C. Fang. Linear programming with inequality constraints via entropic perturbation. *International Journal of Mathematics and Mathematical Sciences*, 19(1):177–184, 1996.
- [25] H. Xu, Z. Rabinovich, S. Dughmi, and M. Tambe. Two-stage security games – exploring information asymmetry. In *The 29th AAAI Conference*, 2015. to appear.
- [26] R. Yang, B. Ford, M. Tambe, and A. Lemieux. Adaptive resource allocation for wildlife protection against illegal poachers. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2014.
- [27] J. Zhuang and V. M. Bier. Reasons for secrecy and deception in Homeland-Security resource allocation. *Risk Analysis*, 30(12):1737–1743, 2010.